Juniper JUNOS

Router Checklist Procedure Guide

02 December 2005

Supplement to the Network Infrastructure Checklist V6R4

DISA
FIELD SECURITY OPERATIONS

This page is intentionally left blank.

TABLE OF CONTENTS

This page is intentionally left blank.

NET0400

Requirement: *The router administrator will ensure neighbor authentication with MD5 is implemented for all routing protocols with all peer routers within the same or between autonomous systems (AS).*

Procedure: Determine what routing protocols have been implemented on the edge with external peers as well as internally. With the exception of external NIPRNet or SIPRNet peers, neighbor authentication must be implemented using MD5. The following routing protocols support MD5: BGP, OSPF, IS-IS, and RIP V2. Following are some sample configurations for BGP and OSPF neighbor authentication.

<u>BGP</u>

```
[edit protocols bgp]
group external-peers {
    type external;
    neighbor 171.69.232.90 {
        peer-as 200;
        authentication-key xxxxx;
    }
    neighbor 171.69.232.100 {
        peer-as 300;
        authentication-key xxxxx;
    }
}
```

Note: The authentication-key statement can be applied at the BGP level, at the group level, or at the neighbor level.

<u>OSPF</u>

```
[edit protocols ospf]

area 0.0.0.0 {
  authentication-type md5
  interface ge-0/0/0.0 {
    authentication-key xxx key-id xxxxxx;
  }
}
```

Note: Authentication has to be enabled for each area. In OSPF, an interface belongs to only one area, which is defined by the interface statement under the *area* statement that it belongs to. The MD5 keyid and password is defined under each interface connected to an OSPF neighbor.

NET0410

Requirement: *The router administrator will restrict BGP connections to known IP addresses of neighbor routers from a trusted AS.*

Procedure: Review the active configuration to ensure that BGP connections are only from known neighbors in a trusted AS by restricting TCP port 179 to specific IP addresses.

Using an Ingress Filter

```
[edit interfaces]
interfaces fe-1/1/1 {
    unit 0 {
        family inet {
            filter {
                input NIPRNet-ingress;
            }
            address 192.168.1.2/32;
        }
    }
}

[edit firewall]
family inet {
    filter NIPRNet-ingress {
        term guard-bgp {
            from {
                source-address {
                    0.0.0.0/0;
                    192.168.1.1/32 except;
                }
                protocol tcp;
                port bgp;
            }
           then {
               syslog;
               discard;
           }
        term  last-accept {
            from {
                destination-address {
                    131.77.5.32/32;
                    131.77.5.61/32;
                }
                protocol tcp;
                destination-port http;
            }
            then accept;
        }
        term default-action {
            then {
               syslog;
               discard;
            }
        }
    }
}
```

Using a Routing Engine Filter

```
[edit interfaces]
lo0 {
    unit 0 {
        family inet {
            filter {
                input protect-routing-engine;
            }
            address 192.168.1.2/32;
        }
    }
}

[edit firewall]
family inet {
   filter protect-routing-engine {
      term guard-bgp {
         from {
            source-address {
               0.0.0.0/0;
               192.168.1.1/32 except;
               }
            protocol tcp;
            port bgp;
         }
         then {
            log;
            discard;
.
.
.
     term default-action {
        then {
         syslog;
        discard;
      }
    }
  }
}
```

NET0430

Requirement: *The IAO/NSO will ensure that an authentication server is used to gain administrative access to all routers.*

Procedure: Verify that an authentication server is required to access the router by reviewing the active configuration. You should find an authentication statement similar to the example below:

```
[edit system]
authentication-order [radius password];
radius-server {
  7.7.7.5 {
    secret xxxxx;
    timeout 20;
  }
}
```

Note: The timeout parameter is amount of time in seconds that the local router waits to receive a response from a RADIUS or TACACS+ server.

NET0440

Requirement: *The IAO/NSO will ensure that when an authentication server is used for administrative access to the router, only one account is defined locally on the router for use in an emergency (i.e., authentication server or connection to the server is down).*

Procedure: Review the active configuration and verify that only one local account has been defined. An example of a local account is shown in the example below:

```
[edit system]
class tier3 {
   idle-timeout 15;
   permissions all;
}
user admin {
   full-name Administrator;
    uid 2000;
    class tier3;
    authentication {
        encrypted-password xxxxxxxxxxx;
    }
}
```

NET0465

Requirement: *The router administrator will ensure that all user accounts are assigned the lowest privilege level that allows them to perform their duties.*

Procedure: When you are using RADIUS or TACACS+ authentication, you can create account templates for authorization purposes that are shared by a group of users. Below is an example configuration with three levels of authorization followed by account templates.

```
[edit system login]
class tier1 {
     idle-timeout 15;
    permissions [configure interface network routing snmp system trace view firewall ];
}
class tier2 {
    idle-timeout 15;
    permissions [admin clear configure interface  network reset routing routing-control
    snmp snmp-control system system-control trace trace-control view maintenance firewall
   firewall-control secret rollback ];
}
class tier3 {
    idle-timeout 15;
    permissions all;
}

/* This is our local superuser account with a local password. */
user admin {
   full-name Administrator;
   uid 2000;
   class tier3;
   authentication {
       encrypted-password xxxxxxx;
   }
}

/* TACACS templates */
user tier1 {
   uid 2001;
   class tier1;
}
user tier2 {
   uid 2002;
   class tier2;
}
user tier3 {
   uid 2003;
```

```
  class tier3;
}
```

Using the example JUNOS configuration above and TACACS configuration below, when a user is using a template account, the CLI username is the login name; however, the privileges, file ownership, and effective user ID are inherited from the template account. The CLI username is sent to the authentication server with the correct password. The server returns the local username (i.e., “tier2”) to the JUNOS software as specified in the authentication server (local-user-name for TACACS+, Juniper-Local-User for RADIUS).

```
user = simon {
.
.
.
.
service = junos-exec {
local-user-name = tier2
allow-commands = "configure"
deny-commands = "shutdown"
}
}
```

Note: *allow-commands* and *deny-commands* override permissions of the class of the template that the local-user-name is associated with.

NET0650

Requirement: *The router administrator will ensure the router console port is configured to time out after 10 minutes or less of inactivity.*

JUNOS procedure: With the exception of root, all user access privileges to a Juniper router are defined in a class. Check the classes that have been defined and examine the idle-timeout parameter. Following is an example:

```
[edit system login]
class superuser-local {
    idle-timeout 10;
    permissions all;
}
```

Notes:

1. There is no default idle-timeout. Without a timeout specified, a login session remains established until a user logs out of the router, even if that session is idle. Unlike IOS, to close idle sessions automatically, you must configure a time limit for each login class.

2. Since the root account does not belong to a class and you can access root via console, disable the ability to login using the root account by making the console insecure.

```
[edit system]
console {
    insecure;
}
```

NET0655

Requirement: *The router administrator will ensure that the router's auxiliary port is disabled.*

JUNOS Procedure: For a Juniper router, the auxiliary port is disabled by default. To enable it, a ports statement in the system hierarchy is required specifying the terminal type. You should **not** find a configuration similar to the following:

```
[edit system]
ports {
    auxiliary {
        type vt100;
    }
}
```

NET0665

Requirement: *The IAO/NSO will ensure that all in-band management connections to the router require passwords.*

Procedure: Any access to a Juniper router requires a login. You can not use CLI unless you are logged in; hence, this will never be a finding.

NET0670

Requirement: *The router administrator will ensure that the router only allows in-band management sessions from authorized IP addresses from the internal network.*

JUNOS Procedure: Review all Juniper router configurations and verify that only authorized internal connections are allowed access to the routing engine via ssh or telnet. Access to the Juniper routing engine is via loopback interface. The configuration should look similar to the following:

```
[edit interfaces]
lo0 {
   unit 0 {
     family inet {
         filter {
             input protect-routing-engine;
         }
         address 192.168.1.2/32;
     }
   }
}


[edit firewall]
family inet {
  filter protect-routing-engine {
    term terminal-access {
       from {
         source-address {
           192.168.1.10;
           192.168.1.11;
          }
        protocol tcp;
        port [ssh telnet];
      }
       then {
          syslog;
          accept;
       }
    }
    term default-action {
       then {
          log;
         discard;
       }
    }
  }
}
```

NET0680

Requirement: The router administrators will ensure in-band management access to the router is restricted to SSH.

Procedure: Review all Juniper router configurations and verify that only access to the routing engine is only via ssh. The configuration should look similar to the following:

```
[edit interfaces]
lo0 {
   unit 0 {
     family inet {
        filter {
          input protect-routing-engine;
         }
         address 192.168.1.2/32;
     }
   }
}


[edit firewall]
family inet {
   filter protect-routing-engine {
       term terminal-access {
           from {
              source-address {
                 192.168.1.10;
                 192.168.1.11;
              }
              protocol tcp;
              port ssh;
           }
           then {
              syslog;
              accept;
           }
       }
.
.
.

       term default-action {
          then {
             syslog;
             discard;
```

```
            }
        }
    }
}
```

NET0685

Requirement: *The router administrator will ensure the timeout for in-band management access is set for no longer than 10 minutes.*

Procedure: With the exception of root, all user access privileges to a Juniper router are defined in a class. Check the classes that have been defined and examine the idle-timeout parameter. Following is an example:

```
[edit system login]
class superuser-local {
            idle-timeout 10;
            permissions all;
        }
```

Note: There is no default idle-timeout; hence, without a timeout specified, a login session remains established until a user logs out of the router, even if that session is idle. Unlike IOS, to close idle sessions automatically, you must configure a time limit for each login class.

When ssh is enabled, all users can use it to access the router---including the root account. This presents two problems:
the root account now be accessed using in-band management
since the root account does not belong to a login class, there is no way to set the idle timeout.

Access to the root account via ssh must be disabled via *root-login deny* command. Following is an example configuration:

```
[edit system]
services {
          ssh {
               root-login deny;
```

NET0690

Requirement: *The router administrator will configure the ACL that is bound to the VTY ports to log permitted and denied access attempts.*

Procedure: Review all Juniper router configurations and verify that all ssh or telnet connection attempts are logged. The configuration should look similar to the following:

```
[edit interfaces]
lo0 {
    unit 0 {
        family inet {
            filter {
                input protect-routing-engine;
             }
             address 192.168.1.2/32;
        }
    }
}
```

```
[edit firewall]
family inet {
   filter protect-routing-engine {
      term terminal-access {
         from {
            source-address {
              192.168.1.10;
              192.168.1.11;
            }
            protocol tcp;
            port [ssh telnet];
         }
         then {
            syslog;
            accept;
         }
      }
      term default-action {
         then {
            syslog;
            discard;
         }
      }
   }
}
```

NET0740

Requirement: *The router administrator will ensure HTTP, FTP, and all BSD r-command servers are disabled.*

Procedure: Under the *edit system services* hierarchy enter a *show* command to verify that the *ftp* and *rlogin* commands are not present. JUNOS does not support the http service.

NET0800

Requirement: *The router administrator will ensure ICMP unreachable notifications, mask replies, and redirects are disabled on all external interfaces of the premise router.*

Procedure: Review the active configuration to determine if controls have been defined to ensure router does not send ICMP unreachables, redirects, and mask replies out any external interfaces.

<u>ICMP Unreachable</u>

1. Protocol Unreachable

The filter used for the routing engine must be configured to silently discard any packets it does not recognize or want. Following would be an example:

```
[edit interfaces]
lo0 {
    unit 0 {
        family inet {
            filter {
                input protect-routing-engine;
            }
            address 192.168.1.2/32;
        }
    }
}

[edit firewall]
family inet {
    filter protect-routing-engine {
        term 1 {
.
.
.
        term default-action {
            then {
                syslog;
                discard;
            }
        }
    }
}
```

2. Host Unreachable

The only method to prevent a Juniper router from sending a Host Unreachable message back to the originator when it receives a packet with a destination address that is not found in its forwarding table, is t define a default route to the *discard* interface. The filter applied to this interface would then silently discard the packets.

```
[edit interfaces]
dsc {
    unit 0 {
        family inet {
            filter {
                input log-discard;
            }
            address 10.1.1.1/32 {
                destination 10.1.1.2;
            }
        }
    }
}

[edit firewall]
family inet {
    filter log-discard {
        term one {
            then {
                 syslog;
                 discard;
            }
        }
    }
}

[edit routing-options]
static {
    route 0.0.0.0/0 next-hop 10.1.1.2 ;
}
```

3. Aggregate and black hole routes

A Juniper router will also send ICMP unreachable messages for packets that have a destination address of an aggregate route as well as a black hole route.

a. Checking aggregate routes

By default, when aggregate routes are installed in a Juniper routing table, the next hop is configured as a *reject* route. Hence the packet is dropped and an ICMP unreachable message is sent to the packet's originator if the aggregate route itself is the result of a routing table longest-match lookup or a packet with a more specific destination under the advertised summary route does not match a more specific route (contributing route). These packets can be quietly dropped by specifying *discard* for an individual route in the *route* part of the *aggregate* statement, or specifying *reject* when you configure the defaults for aggregate routes.

```
[edit routing-options]
aggregate {
    route 192.168.0.0/17 discard ;
```

or

```
[edit routing-options]
aggregate {
    defaults {
          active;
          discard;
          community 2:333;
    }
}
```

Note: You can also issue the operational command *show route protocol aggregate* to determine if *discard* or *reject* option is used.

b. Checking black hole routes

```
[edit routing-options]
static {
   route 0.0.0.0/8 discard;
   route 1.0.0.0/8 discard;
   route 5.0.0.0/8 discard;
.
```

ICMP Redirects

Under the *edit system* hierarchy enter a *show* command to verify that the *no-redirects* command is present on all Juniper routers. This restriction can also be enforced by including the *no-redirects* statement under each active interface.

```
[edit system]
no-redirects;
```

or

```
[edit interfaces]
fe-2/0/1 {
       description "NIPRNet link";
       unit 0 {
          family inet {
                no-redirects;
                filter {
                    input ingress-filter;
                }
                address 121.70.11.68/29;
            }
        }
    }
}
```

ICMP Mask Reply

JUNOS has no option to not reply to an ICMP Mask Request message. Consequently, to ensure that the router does not send any ICMP Mask Reply messages in response to a mask request, include a term statement in the routing engine firewall to drop any masks requests sent to it.

```
[edit interfaces]
lo0 {
    unit 0 {
        family inet {
            filter {
                input protect-routing-engine;
            }
            address 192.168.1.2/32;
        }
    }
}

[edit firewall]
family inet {
    filter protect-routing-engine {
        term icmp-mask-request {
            from {
                protocol icmp;
                icmp-type mask-request;
            }
            then {
                log;
                discard;
            }
        }
    }
}
```

NET0810

Requirement: *The IAO/NSO will ensure that two Network Time Protocol (NTP) servers are defined on the premise router to synchronize its time.*

Procedure: Review the router configurations and verify that NTP servers have been defined similar to the following example:

```
[edit system]
ntp {
   boot-server 129.237.32.2;
   server 129.237.32.2;
   server 142.181.31.6;
}
```

Note: The *boot-server* statement identifies the server from which the initial time of day and date is obtained when the router boots. The *server* statements identifies the NTP servers used for periodic time synchronization.

NET0820

Requirement: *The IAO/NSO will ensure that the DNS servers are defined if the router is configured as a client resolver.*

Procedure: Review the active configuration to ensure that DNS servers have been defined similar to the following example:

```
[edit system]
name server {
   192.168.1.253;
   192.168.1.254;
}
```

Note: Since JUNOS will not send a DNS query to resolve names to IP addresses if a name server is not defined—unlike Cisco that will send a broadcast—this will never be a finding.

NET0890

Requirement: *The router administrator will restrict SNMP access to the router from only authorized internal IP addresses.*

Procedure: Review all Juniper routers to ensure that SNMP access is limited to specific IP hosts using a configuration similar to the following:
.

```
snmp {
  interface ge-0/1/0.0;
  community xxxxxxxxx {
    authorization read-only;
    clients {
      default restrict;
      7.7.7.5/32;
    }
  }
}
```

Note: if the *clients* statement is not present, then all clients are allowed.

NET0894

Requirement: *The router administrator will ensure SNMP is only enabled in the read mode; Read/Write is not enabled unless approved and documented by the IAO/NSO.*

Procedure: Review all Juniper routers to ensure that SNMP access is limited to specific IP hosts using a configuration similar to the following:

```
snmp {
  interface ge-0/1/0.0;
  community xxxxxxxxx {
    authorization read-only;
    clients {
      default restrict;
      7.7.7.5/30;
    }
  }
}
```

NET0910

Requirement: *The router administrator will utilize ingress and egress ACLs to restrict traffic in accordance with the guidelines contained in DOD Instruction 8551.1 for all ports and protocols required for operational commitments.*

Procedure:

1. Determine Boundary- Determined by Connectivity, not Destination. ACLs use source and destination addresses. PPS defines boundary by physical connectivity. All Federal Agencies are not DOD!

2. NOTE on Enclave to Enclave: If data traffic between Enclaves transverses a router not owned by the Enclave's DAA then it falls into the "Boundary 7&8 DoD Network to Enclave or other applicable category for that particular connectivity.

3. Block by specifying ports on permit statements when in deny-by-default **<u>or</u>** explicitly block all known red ports.

4. All ports and protocols allowed into the enclave should be registered in the PPS database.

Review the premise router configuration to ensure ACLs are in place to restrict inbound IP addresses are filtered to permit only green or yellow ports. Red and yellow ports are permitted with conditions noted on the Category Assignment List (CAL). A DSAWG 2 year expiration date listed on the PPS CAL will indicate expiration of permits for particular red ports. The router configuration should look <u>similar</u> to following highlighted:

```
[edit interfaces]
fe-2/0/10 {
        description "to NIPRNet core router";
        speed 100m;
        link-mode full-duplex;
        unit 0 {
            family inet {
                filter {
                    input ingress-filter;
                }
                address 199.36.92.1/30;
            }
        }
    }

[edit firewall]
family inet {
    filter ingress-filter {
        term term-1 {
            from {
```

```
           source-address {
              <internal network range>/<prefix>;
              0.0.0.0/8;
              10.0.0.0/8;
              127.0.0.0/8;
              169.254.0.0/16;
              172.16.0.0/12;
              192.0.2.0/24;
              192.168.0.0/16;
              224.0.0.0/4;
              240.0.0.0/5;
           }
        }
        then discard;
     }
   term pps-access {
      from {
        source-address {
          192.168.1.10;
          192.168.1.11;
         }
       protocol tcp;
       port [ssh telnet];
     }
      then {
         syslog;
         accept;
      }
   }
 }
```

NET0920

Requirement: *The router administrator will bind the ingress ACL filtering packets entering the network to the external interface, and bind the egress ACL filtering packets leaving the network to the internal interface—both on an inbound direction.*

Note: All filters must be applied to the appropriate interfaces on an inbound direction. Ingress filtering is applied to all traffic entering the enclave; hence, this filter would be bound to all external interfaces. Since egress filtering is applied to all traffic leaving the enclave, this filter would be bound to all internal interfaces.

Procedure: Review the active configuration of the premise router and verify that all interfaces, with the exception of the loopback, fxp0, and fxp1 interfaces, have the appropriate ingress or egress firewall filter applied with an inbound direction. An example configuration is depicted below:

```
[edit interfaces]
interfaces fe-1/1/1 {
   unit 0 {
      family inet {
         filter {
            input NIPRNet-ingress;
         }
         address 199.36.92.1/32;
      }
   }
}
```

NET0940

Requirement: *The router administrators will restrict the premise router from accepting any inbound IP packets with a source address that contain an IP address from the internal network, any local host loop back address (127.0.0.0/8), the link-local IP address range (169.254.0.0/16), or any reserved private addresses in the source field.*

Procedure: Review the Juniper premise router configuration to ensure firewall filters are in place to restrict inbound IP addresses. The ingress firewall filter should look similar to following:

```
[edit interfaces]
fe-2/0/10 {
    description "to NIPRNet core router";
    speed 100m;
    link-mode full-duplex;
    unit 0 {
      family inet {
        filter {
          input ingress-filter;
        }
        address 199.36.92.1/30;
      }
    }
  }
```

```
[edit firewall]
family inet {
  filter ingress-filter {
     term term-1 {
       from {
         source-address {
           <internal network range>/<prefix>;
           0.0.0.0/8;
           10.0.0.0/8;
           127.0.0.0/8;
           169.254.0.0/16;
           172.16.0.0/12;
           192.0.2.0/24;
           192.168.0.0/16;
           224.0.0.0/4;
           240.0.0.0/5;
         }
       }
       then discard;
     }
   }
}
```

NET0950

Requirement: *The router administrator will restrict the router from accepting any outbound IP packet that contains an illegitimate address in the source address field via egress ACL or by enabling Unicast Reverse Path Forwarding.*

Procedure for egress ACL: Review the premise router configuration to ensure egress filters are in place on all internal interfaces to restrict the router from accepting outbound IP packets that contain an external IP address in the source field. In order to comply with the deny-by-default policy, permit statements for those ports that are allowed will have to be defined followed by the deny any statement. The permit statements must qualify the source address with the internal network address range. Following is an example:

```
[edit interfaces]
fe-2/0/10 {
    description "downstream link to our network”;
    unit 0 {
      family inet {
        filter {
          input egress-filter;
        }
        address 199.36.90.1/24;
      }
    }
  }
}

[edit firewall]
family inet {
  filter egress-filter {
    term term-established {
      from {
        protocol tcp;
        tcp-established;
      }
      then accept;
    }
    term ext-DNS {
      from {
        source-address {
           201.111.2.130/32;
         }
         protocol udp;
         source-port 53;
        }
        then accept;
```

```
         }
      }
      term http-ftp {
         from {
            source-address {
               201.111.2.0/24;
            }
            protocol tcp;
            destination-port [20 21 80 443 ];
            }
            then accept;
         }
      }
.
.
.

      term default-action {
         then {
             syslog;
             discard;
         }
      }
   }
}
```

Procedure for Unicast Reverse Path Forwarding: Review the premise router configuration to ensure RPF has been configured on all internal interfaces. Following is an example configuration:

```
[edit interfaces]
fe-2/0/10 {
      description "downstream link to our network”;
      unit 0 {
         family inet {
             rpf-check fail-filter filter-log-rpf-failure;
            filter {
               input egress-filter;
            }
            address 199.36.90.1/24;
         }
      }
   }
}

[edit firewall]
family inet {
    filter filter-log-rpf-failure {
         term log-and-drop {
               then {
                   log;
                   discard;
               }
         }
    }
}
```

Note: To consider only active paths during the unicast RPF check, include the *active-paths* option. To consider all feasible paths during the unicast RPF check, include the *feasible-paths* option.

A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router forwards the packet to the destination address. If it is not from a valid path, the router discards the packet.

Best practice is to implement Unicast RPF downstream from the premise router, preferably at the distribution layer or at the edges of the network. The further downstream Unicast RPF is applied, the finer the granularity you have in mitigating address spoofing and in identifying the sources of spoofed addresses. Applying Unicast RPF on the premise router helps mitigate attacks from many downstream networks and it is easier to administer, but it makes it harder to identify the source of the attack.

NET0960

Requirement: *The IAO/NSO will implement features provided by the router to protect servers from any TCP SYN flood attacks from an outside network.*

JUNOS does not have a method to proxy or watch over TCP connection attempts. However, it does have rate limiting mechanisms that can be used to mitigate a SYN flood attack against a network or targeted hosts. Rate limiting TCP SYN packets with JUNOS can be performed by including rate-limiting on the ingress firewall filter assigned to external-facing interfaces. Rate limiting can be confiugred to limit the amount of bandwidth consumed as well as the maximum burst size of the TCP SYN traffic. The configuration should look similar to the following:

```
interfaces {
   fe-2/0/10 {
      description "to NIPRNet core router";
      speed 100m;
      link-mode full-duplex;
      unit 0 {
         family inet {
            filter {
               input ingress-filter;
            }
            address 199.36.92.1/30;
         }
      }
   }

firewall {
   filter ingress-filter {
      policer tcp-syn-control {
         if-exceeding {
            bandwidth-limit 5000k;
            burst-size-limit 150k;
         }
         then discard;
      }

/* Rate limit TCP control traffic from external sources */

      term 1 {
         from {
            protocol tcp;
            tcp-flags (syn & !ack) | fin | rst";
         }
         then {
            policer tcp-syn-control;
```

```
                accept;
            }
        }
    }
}
```

NET0965

Requirement: *The router administrator will set the maximum wait interval for establishing a TCP connection request to the router to 10 seconds or less, or implement a feature to rate-limit TCP SYN traffic destined to the router.*

JUNOS Procedure: Review the premise or edge router confiugration to verify that one of the following methods have been implmented to rate limit connection attempts to the router.

Rate limiting by Services

```
services {
   ssh {
      root-login deny;
      protocol-version v2;
      connection-limit <limi>t;
      rate-limit <limit>;
   }
   telnet {
      connection-limit <limit>;
      rate-limit <limit>;
   }
   ftp {
      connection-limit <limit>;
      rate-limit <limit>;
   }
```

Rate limiting TCP SYN traffic to Routing Engine

```
interfaces {
   lo0 {
      unit 0 {
         family inet {
            filter {
               input protect-routing-engine;
            }
            address 192.168.1.2/32;
         }
      }
   }
}
```

```
policy-options {
/* Addresses to be used in protect-routing-engine filter */
   prefix-list ssh-connect {
      192.168.1.4/32;
      192.168.1.5/32;
   }
   prefix-list bgp-connect {
      5.5.5.1/32;
   }
}

firewall {
   filter protect-routing-engine {
      policer tcp-syn-control {
         if-exceeding {
            bandwidth-limit 500k;
            burst-size-limit 15k;
         }
         then discard;
      }

/* Rate limit TCP control traffic from trusted sources */

      term 1 {
         from {
            source-prefix-list {
               ssh-connect;
               bgp-connect;
            }
            protocol tcp;
            tcp-flags (syn & !ack) | fin | rst";
         }
         then {
            policer tcp-syn-control;
            accept;
         }
      }
   }
}
```

NET0980

Requirement: *The router administrator will block all inbound ICMP messages with the exception of Echo Reply (type 0), and Time Exceeded (type 11). ICMP message number 3, code 4, are permitted inbound with the following exception: Must be denied from external AG addresses, otherwise permitted.*

Procedure: Review the Juniper premise router configuration to ensure that the ingress filter blocks all inbound ICMP traffic message types with the exception of Echo Reply (type 0), Time Exceeded (type 11), and Destination Unreachable (type 3). The filter should look similar to the following:

```
[edit firewall]
family inet {
    filter ingress-filter {
        term icmp-in-good {
            from {
                protocol icmp;
                icmp-type [echo-reply time-exceeded unreachable];
            }
            then {
                accept;
            }
        }
        term icmp-in-bad {
            from {
                protocol icmp;
            }
            then {
                syslog;
                discard;
            }
         }
     }
}
```

NET0990

Requirement: *The router administrator will block outbound ICMP traffic message types except Echo Request (type 8), Parameter Problem (type 12), and Source Quench (type 4) Destination Unreachable - Fragmentation Needed and Don't Fragment was Set (type3, code 4).*

Procedure: Review the Juniper premise router configuration to ensure that the egress filter is bound to the proper interfaces to block outbound ICMP traffic message types except Echo Request, Parameter Problem, Source Quench, and Destination Unreachable - Fragmentation Needed and Don't Fragment was Set. The configuration should look similar to the following:

```
[edit interfaces]
fe-2/0/10 {
      description "link to our network”;
      unit 0 {
         family inet {
            filter {
               input egress-filter;
            }
            address 199.36.90.1/24;
         }
      }
   }
}

[edit firewall]
family inet {
    filter egress-filter {
        term icmp-out-good {
            from {
                source-address {
                    199.36.90.0/24;
                }
                protocol icmp;
                icmp-type [ echo-request source-quench parameter-problem ];
            }
            then {
                accept;
            }
        }
        term icmp-PMTU-D {
            from {
                source-address {
                    199.36.90.0/24;
                }
                protocol icmp;
                icmp-type [ unreachable ];
                icmp-code [ fragmentation-needed ];
            }
            then {
                accept;
            }
        }
        term icmp-out-bad {
            from {
                protocol icmp;
            }
            then {
                syslog;
                discard;
            }
         }
    }
}
```

NET1000

Requirement: *The router administrators will block all inbound traceroutes to prevent network discovery by unauthorized users.*

Procedure: Review the premise router configuration to ensure that an ingress filter is in place to block inbound UDP 33400 through 34400 as well as any packet with the value of 82 in the IP Options field.

```
[edit interfaces]
fe-2/0/10 {
    description "NIPRNet link”;
    unit 0 {
      family inet {
        filter {
          input ingress-filter;
        }
        address 201.126.90.1/30;
      }
    }
  }
}

[edit firewall]
family inet {
   filter ingress-filter {
     term one {

.
.
.
.
     term block-old-traceroute {
        from {
          protocol udp;
          destination-port [ 33400-34400 ];
        }
        then {
          syslog;
          discard;
      }
    }
     term block-new-traceroute {
        from {
          ip-options 82;
        }
```

```
            then {
                syslog;
                discard;
          }
       }
    }
}
```

Note: Resource Reservation Protocol (RSVP) used by MPLS, Internet Group Management Protocol Version 2 (IGMPv2), and other protocols that use the IP options field may not function.

NET1020

Requirement: *The router administrator will ensure that all attempts to any port, protocol, or service that is denied are logged.*

Procedure: Review the active configuration of the premise router and verify that both the router's ingress and egress firewall filters have a *log* or *syslog* statement for every *discard* or *reject* statement. An example configuration is depicted below:

```
[edit firewall]
family inet {
  filter NIPRNet-ingress {
    term first-accept {
      from {
.
.
.
        }
        then accept;
    }
    term  last-accept {
      from {
        destination-address {
          131.77.5.32/32;
          131.77.5.61/32;
         }
         protocol tcp;
         destination-port http;
      }
      then accept;
    }
   term default-action {
     then {
       syslog;
       discard;
     }
   }
  }
}
```

Note: The *log* action command will satisfy this requirement. However, the *syslog* action command will be required to satisfy a subsequent requirement for logging all data to a syslog server. JUNOS will only log to the buffer if the *log* action command is used.

NET1021

Requirement: *The router administrator will configure all routers to log severity levels 0 through 6 and send log data to a syslog server.*

Procedure: Review all router configurations to ensure that all routers log messages for severity levels 0 through 6. By specifying *info*, all severity levels above will be included.

Logging Level	**Severity Level**	**Description**
Emergencies	0	
Alerts	1	Immediate Action Required
Critical	2	Critical Conditions
Errors	3	Error Conditions
Warnings	4	Warning Conditions
Notifications	5	Normal but Significant Conditions
Informational	6	Informational Messages
Debugging	7	Debugging Messages

A sample configuration would look similar to the following:

```
[edit system syslog]
syslog {
    host 192.168.1.22 {
        any info;
        facility-override local7;
    }
}
```

NET1070

Requirement: *The IAO/NSO will authorize and maintain justification for all tftp implementations.*

Procedure: Configuration files can be copied to and from the router using the *file copy* command in operational mode or *save* command while in configuration mode. The destination address is specified on the command line—never preconfigured. Destinations can be the router’s flash (path/filename), hard drive (/var/filename), removable media (a:filename), FTP server (ftp://hostname/path/filename), TFTP server (tftp://hostname/path/filename), HTTP server (http://hostname/path/file), or an Secure Copy Protocol (SCP) client (scp://hostname/path/filename).

Unless TFTP, FTP, or HTTP is specified in the command string, both the *save* and *file copy* commands will utilize Secure Copy Protocol, which uses the SSH authentication and encryption framework, to securely copy files to and from a remote host. Interview the router administrator to determine what method is used. If the site uses TFTP or HTTP with the *save* or *file copy* command, this is a finding.